PhD in Information Technology and Electrical Engineering

Università degli Studi di Napoli Federico II

PhD Student: Antonio Montieri

XXXII Cycle

Training and Research Activities Report – Third Year

Tutor: Prof. Antonio Pescapè

Information

I am **Antonio Montieri** and I received a M.Sc. degree cum laude in Computer Engineering from the University of Napoli Federico II in July 2015. Currently, I am a PhD Candidate of the XXXII Cycle of the Information Technology and Electrical Engineering (ITEE) PhD program at the Department of Electrical Engineering and Information Technology (DIETI) of the University of Napoli Federico II. My fellowship is financed by a university ITEE grant. My tutor is prof. **Antonio Pescapè** and I am a member of the TRAFFIC research group, part of the larger COMICS, whose activities are carried out in the field of Computer Networks.

Study and Training activities

During the third year of PhD program, I attended the courses and seminars reported in the following. In June 2019, I also attended the 9th TMA PhD School on Traffic Measurement & Analysis at the Conservatoire National des Arts et Métiers (Cnam) in Paris (French) that has been included in the list of seminars.

Courses

1. *Internet censorship: enforcement, detection, and circumvention*, ad hoc module, Prof. Giuseppe Aceto, 07/05/2019 – 09/05/2019, 2 Credits.
2. *Accelerated Computing with Cuda C/C++*, ad hoc module, Prof. Luigi Troiano, 25/11/2019, 0.5 Credits.
3. *Deep Learning for Computer Vision*, ad hoc module, Prof. Luigi Troiano, 16/12/2019, 0.5 Credits.

Seminars

1. *Applying Semi-Supervised Learning to App Store Analysis*, Dr. Daniel Rodriguez, 12/07/2019, 0.2 Credits.
2. *PhD School: 9th TMA PhD School on Traffic Measurement & Analysis*, Sara Dickinson; Narseo Vallina-Rodriguez; Mirja Kühlewind; Tim Bruijnzeels, 17/06/2019 – 18/06/2019 (15 hours), 1.5 Credits.

Credits Summary

Finally, I provide a table reporting a summary of the credits obtained attending modules and seminars and doing research activities.

	Credits year 1								Credits year 2								Credits year 3								
	Estimated	1 bimonth	2 bimonth	3 bimonth	4 bimonth	5 bimonth	6 bimonth	Summary	Estimated	1 bimonth	2 bimonth	3 bimonth	4 bimonth	5 bimonth	6 bimonth	Summary	Estimated	1 bimonth	2 bimonth	3 bimonth	4 bimonth	5 bimonth	6 bimonth	Summary	Total
Modules	20	4	3	0	6	4	6	**23**	10	0	3	2,4	0	1,2	0	**6,6**	5	0	2	0	0	0,5	0,5	**3**	**32,6**
Seminars	10	4,9	1,8	2	0,5	0,3	0,3	**9,8**	5	0,8	0,4	3,9	0	0,4	0	**5,5**	0	0	0	1,7	0	0	0	**1,7**	**17**
Research	30	1,1	5,2	8	3,5	5,7	6,7	**30,2**	45	9,2	6,6	3,7	10	8,4	10	**47,9**	55	9,5	7,5	8,3	9	9	9	**52,3**	**130,4**
	60	10	10	10	10	10	13	**63**	60	10	10	10	10	10	10	**60**	60	9,5	9,5	10	9	9,5	9,5	**57**	**180**

Research Activity

During my PhD, I have worked (and I am currently working) in the context of monitoring and management of computer networks, with specific focus on mobile and encrypted traffic classification. Specifically, with Traffic Classification (TC) we refer to the process of associating network traffic with specific applications generating it. This research activity aims at shedding light on the changing nature of the traffic generated by smartphones (and other handheld devices) whose deep usage in everyday life is growing more and more. Indeed, various tools and middle-boxes, such as security/quality-of-service enforcement devices and network monitors, rely on the knowledge of the application generating the traffic and thus are limited (or impaired) when this requirement is not completely satisfied. Hence, the research activity carried out in my third year of PhD, can be summarized as follows.

Techniques for Mobile and Encrypted Traffic Classification via Deep Learning

I have deepened the research on (advanced) Deep Learning (DL)-based TC techniques, which allows to train classifiers directly from input data by automatically learning structured and complex feature representations, overcoming the limitations of "traditional" Machine Learning (ML) classifiers based on handcrafted (domain-expert driven) features. In more detail, I have explored and implemented multi-modal architectures that are able to exploit the intrinsic heterogeneous nature of traffic data and can capitalize the different views (viz. modalities) of the same traffic object (e.g., raw payload or header fields). Also, I have started the research on multi-task architectures being in charge of providing inference for TC problems (e.g., inferring both the traffic-type and the specific application).

I have also tried to face a (possible) limitation of DL-based TC, that is the generation of learning networks with very dense and complex structure, whose training might be excessively slow and computational demanding with respect to the timeliness and computational constraints of network domain. This issue is even more urgent in the fast-evolving mobile scenario. With this aim, I have investigated the integration of the Big Data (BD) framework with DL-based TC architectures. Indeed, BD solutions provide processing frameworks that can parallelize the classification task by splitting the network data and distributing it across different workers cooperating under the coordination of a single master.

Hierarchical Approaches for the Classification of Anonymity Tools' Traffic

I have applied hierarchical ML-based classification scheme to the classification at different levels of anonymous traffic (i.e. the traffic generated by means of anonymity tools such as Tor, I2P, and JonDonym). Hierarchical classification allows fine-grained tuning and design of classifiers, potentially leading to performance gains, and it also brings a number of "practical" benefits by design, at cost of moderate complexity increase.

Mobile Traffic Analysis in Large-Scale Scenarios

This research activity is made in collaboration with the IMDEA Networks Institute, Leganes (Madrid), Spain. It regards the characterization and classification of mobile traffic in large-scale scenarios and from multiple vantage points, and the investigation of mobile-traffic privacy and 3rd-party tracking ecosystem.

Research Description

Techniques for mobile and encrypted traffic classification via Deep Learning

The efficacy of security and quality-of-service enforcement devices, as well as network monitors, is limited (or qualitatively hampered) when there is no accurate knowledge of the application generating the traffic. The process inferring such information, known as network Traffic Classification (TC), has a long-standing application in many fields [1] and is facing unprecedented challenges due to the users' massive shift toward mobile devices (as witnessed by recent Internet traffic evaluations [2]), leading to a multi-faceted and evolving composition of network traffic [3].

Hence, the appeal of mobile TC has bloomed nowadays, nurtured (other than usual TC drivers, e.g. service differentiation) by valuable profiling information (e.g., to advertisers, security agencies, and insurance companies), while also implying privacy downsides (e.g., recognition of context-sensitive applications, such as dating and health ones, and bring-your-own-devices policies). Equally important, the effort towards the protection of privacy and security has fueled the widespread adoption [4, 5] of encrypted protocols (TLS). This shift, together with the use of dynamic transport ports or the clustering on a few well-known (and commonly unblocked) ports, resulted in the hampering of accurate TC, as both Packet Inspection (DPI) and port-based techniques become ineffective [1]. Moreover, other than the ET issue, mobile TC comes with exacerbated challenges and requirements due to a large number of apps to discriminate from and the automatic frequent updates of the apps—leading to inadequate number of training samples per app and hindering the achievement of targeted performance.

In this context, ML classifiers have proved to be a good fit, since they suit also Encrypted Traffic (ET) while not expressly relying on port information [6, 7, 8]. However, their usual form resorts to the process of obtaining handcrafted (domain-expert driven) features (e.g., packet sequence statistics), which is time-consuming, unsuited to automation, and it is becoming rapidly outdated when compared to the evolution and mix of mobile traffic, being a constantly moving target, and precluding the design of accurate and up-to-date mobile-traffic classifiers [8, 9] with "traditional" ML approaches. Therefore, DL is emerging as the stepping stone toward the fulfillment of high performance in the dynamic and challenging (encrypted) TC contexts, allowing to train classifiers directly from input data by automatically distilling structured (and complex) feature representations [10].

Furthermore, most of these DL-based efforts have focused on one type of input information (e.g., payload bytes or header fields), despite traffic data being naturally "multi-modal". Indeed, the main asset of multi-modal DL is the ability of automatically learning a hierarchical representation exploiting jointly all the available modalities, instead of handcrafting modality-specific features for a given ML approach [17].

The challenges concerning the higher complexity (and corresponding training time) of DL architectures as compared to ML-based classifiers, and the lack of accurately-labeled traffic datasets are also taken into account.

This activity has led to the publication of two journal papers [J1, J2] and two conference papers [C1, C2], made in collaboration with the other members of the TRAFFIC research group. Additionally, we have submitted the journal paper [J4], currently under the first round of review.

Specifically, in [J1], we propose the design of mobile traffic classifiers (able to operate with ET) via the adoption of the DL umbrella. To this end, this work resorts to the development of a systematic framework for the design of novel DL-based TC architectures and comparison of existing ones. We apply this framework for a critical analysis of several non-mobile-specific DL classifiers recently appeared in TC literature [11, 12, 13, 14, 15, 16]. In detail, the proposed framework dissects the DL-based TC problem from different viewpoints: (A) the TC object adopted, (B) the type (and the amount) of input data fed to the DL classifier, (C) the DL architecture employed, and (D) the required set of performance measures for an objective and comprehensive evaluation. The outcomes of this work underline the deficiencies of current DL-based traffic classifiers and the need for: (i) unbiased, informative, and heterogeneous inputs extrapolated from traffic data, (ii) sophisticated DL architectures, and (iii) a rigorous and multifaceted performance evaluation. Indeed, this study represents a first attempt to address (i) and (ii) issues, being also a "safe" groundwork for paving the way to the design of accurate DL-based classifiers coping with highly-diverse mobile traffic, whereas it provides designers with a fine-level performance evaluation workbench (iii).

Moving beyond, in [J2], to address the challenge of taking advantage of the heterogeneous nature of network traffic data, we propose a novel MultImodal DL-based MobilE TraffIc Classification (MIMETIC) framework, having the capability of exploiting effectively the different views of a traffic object, by capturing both intra- and inter-modalities dependence. Results, based on three datasets collected by human users, highlight a performance improvement (in terms of both concise and fine-grained measures) while reporting a lower training time (more than three times) with respect to existing (single-modality) DL-based traffic classifiers.

Starting from the expertise gained with the above-mentioned studies, in [J4] (currently under the first round of review), we first give an overview of the key network traffic analysis problems where DL is foreseen as attractive, due to their common capitalization of network-level data in automatic fashion. Secondly, we categorize the state-of-the-art in DL-based TC toward its effective application in mobile and encrypted context. To pinpoint and overcome the limitations found in literature, we propose a general framework for DL-based mobile and encrypted TC, based on a rigorous definition of its milestones: the choice of the traffic object, the definition of the input(s), the number of (simultaneous) TC tasks required and the corresponding DL architecture. We validate two implementations of our framework on three recent mobile traffic datasets, also surfacing future directions toward sophisticated DL-based mobile TC.

While DL represents a promising solution toward high performance and reduced domain expertise in the design of accurate mobile traffic classifiers, it results in higher completion times, in turn suggesting the application of the Big-Data (BD) paradigm. In [C1], we investigate for the first time BD-enabled classification of encrypted mobile traffic using DL from a general standpoint, (a) defining general design guidelines, (b) leveraging a public-cloud platform, and (c) resorting to a realistic experimental setup. We found that, while BD represents a transparent accelerator for some tasks, this is not the case for the training phase of DL architectures for TC, requiring a specific BD-informed design. The experimental setup is built upon a three-dimensional investigation path in the BD adoption, namely: (i) completion time, (ii) deployment costs, and (iii) classification performance, highlighting relevant non-trivial trade-offs.

Finally, in [C2], we introduce and describe MIRAGE, a reproducible architecture for mobile-app traffic capture and ground-truth creation. MIRAGE allows to support data-driven TC approaches that require a reliably-labeled dataset (viz. with a reliable ground-truth), possibly human-generated, to ensure their proper design, realization, and evaluation. Indeed, the existing design solutions are mainly evaluated on private traffic traces, and only a few public datasets are available, thus clearly limiting repeatability and further advances on the topic. To fill this gap, the outcome of this system is a human-generated dataset for mobile traffic analysis (with associated ground-truth) having the goal of advancing the state-of-the-art in mobile app traffic analysis. The dataset has been collected by more than 200 human users on a volunteer basis and contains the traffic generated by more than 50 Android apps. To the best of our knowledge, no dataset with such characteristics has been released to the research community to date. In [C2], we also provide a first statistical characterization of the mobile-app traffic in this dataset.

Hierarchical Approaches for the Classification of Anonymity Tools' Traffic

I have also applied ML-based techniques for the classification of the traffic generated by Anonymity Tools (ATs). Specifically, among the several ATs developed in recent years, I have considered the Onion Router (Tor) [18], the Invisible Internet Project (I2P) [19], and JonDonym [20] being the most popular ones. Starting from the work [22] published during the second year of my PhD programme, in collaboration with other members of the research group, we have published the journal paper [J3] in which we propose TC of anonymity tools (and deeper, of their running services and applications) via a truly hierarchical approach. Capitalizing the Anon17 public dataset [21] released in 2017 containing anonymity traffic, we provide an in-depth analysis of TC and we compare the proposed hierarchical approach with a flat counterpart [22]. The proposed framework is investigated in both the usual TC setup and its "early" variant (i.e. only the first segments of traffic aggregate are used to take a decision). Results highlight a general improvement over the flat approach in terms of all the classification metrics. Moreover, fine-grain performance investigation allows to (a) demonstrate lower severity of errors incurred by the hierarchical approach (as opposed to the flat case) and (b) highlight poorly-classifiable services/applications of each anonymity tool, gathering useful feedback on their privacy-level.

Mobile Traffic Analysis in Large-Scale Scenarios

The present research activity is made in collaboration with the IMDEA Networks Institute, Leganes (Madrid), Spain. It has led to the preparation of the journal paper [J5], we plan to submit on February 2020, made in collaboration also with the International Computer Science Institute (ICSI) and the University of Massachusetts. In this paper, we aim to study and characterize mobile traffic using anonymized traffic traces collected by Lumen, a mobile privacy enhancing and traffic analysis app on Android. We provide a comprehensive view of the world of mobile traffic by studying the diverse and often unique ways in which apps communicate over the network, and use the insights gained from our characterization to extract features to train and evaluate a range of traffic classification methods operating at different simulated vantage points on the path, from ISPs to app servers and CDNs, to ascribe traffic flows to their originating apps. Our characterization of mobile traffic reveals that various factors such as developer decisions, OS environment, and user interactions all have a hand in how an app behaves, and that a significant amount of traffic traces and app fingerprints are needed to efficiently and robustly classify mobile traffic at scale.

Finally, we demonstrate how mobile traffic classification is computationally costly, especially in open-world scenarios and with higher scale.

Collaborations

We strictly collaborate with **Huawei Technologies Co. Ltd.** in the context of the research activity related to mobile and encrypted traffic classification.

Furthermore, I am collaborating with the **IMDEA Networks Institute**, Leganes (Madrid), Spain, on the research topics defined during the periods of study and research abroad detailed in the following appropriate section.

Products

Journal Papers

1. [J1] Giuseppe Aceto, Domenico Ciuonzo, **Antonio Montieri**, Antonio Pescapè, *Mobile Encrypted Traffic Classification using Deep Learning: Experimental Evaluation, Lessons Learned, and Challenges*, IEEE Transactions on Network and Service Management (TNSM), Volume 16, Issue 2, June 2019.
2. [J2] Giuseppe Aceto, Domenico Ciuonzo, **Antonio Montieri**, Antonio Pescapè, *MIMETIC: Mobile Encrypted Traffic Classification using Multimodal Deep Learning*, Elsevier Computer Networks (COMNET), Volume 165, 24 December 2019.
3. [J3] **Antonio Montieri**, Domenico Ciuonzo, Giampaolo Bovenzi, Giuseppe Aceto, Valerio Persico, Antonio Pescapè, *A Dive into the Dark Web: Hierarchical Traffic Classification of Anonymity Tools*, IEEE Transactions on Network Science and Engineering (TNSE), Early Access, 2019.

Conference Papers

1. [C1] Giuseppe Aceto, Domenico Ciuonzo, **Antonio Montieri**, Valerio Persico, Antonio Pescapè, *Know your Big Data Trade-offs when Classifying Encrypted Mobile Traffic with Deep Learning*, 3rd Network Traffic Measurement and Analysis Conference (TMA 2019), June 17-21, 2019, Paris, France.
2. [C2] Giuseppe Aceto, Domenico Ciuonzo, **Antonio Montieri**, Valerio Persico, Antonio Pescapè, *MIRAGE: Mobile-app Traffic Capture and Ground-truth Creation*, 4th IEEE International Conference on Computing, Communications and Security (ICCCS 2019), October 10-12, 2019, Rome, Italy. ***Best Paper Award ICCCS 2019.***

Papers Under Review

1. [J4] Giuseppe Aceto, Domenico Ciuonzo, **Antonio Montieri**, Antonio Pescapè, *Toward Effective Mobile Encrypted Traffic Classification through Deep Learning*, Elsevier Future Generation Computer Systems (FGCS), under the first round of review, 2020.

Papers in Preparation

2. [J5] Abbas Razaghpanah, Álvaro Feal, Julien Gamba, Narseo Vallina-Rodriguez, **Antonio Montieri**, Antonio Pescapè, Mark Allman, Philippa Gill, Transactions on Networking (TON), expected submission February 2020 (in collaboration with the IMDEA Networks Institute, the International Computer Science Institute, and the University of Massachusetts).

Conferences and Seminars

I attended the following conferences:

1. *The 2019 Network Traffic Measurement and Analysis Conference* (TMA 2019), June 17-21, 2019, Paris, France.

I also made the following presentations:

1. *Know your Big Data Trade-offs when Classifying Encrypted Mobile Traffic with Deep Learning*, presented at the 2019 Network Traffic Measurement and Analysis Conference (TMA 2019), June 17-21, 2019, Paris, France.

Activity Abroad

During my third PhD year I have carried out a period of study and research abroad **from 24/01/2019 to 10/04/2019** at the **IMDEA Networks Institute**, Leganes (Madrid), Spain, under the supervision of the Prof. Narseo Vallina-Rodriguez. Study and research activity has concerned the analysis of mobile traffic in large-scale scenarios, and the investigation of mobile-traffic privacy and the 3rd-party tracking ecosystem.

Tutorship

Teaching assistant at the B.Sc. course of **Reti di Calcolatori I** and M.Sc. course of **Analisi e Prestazioni di Internet**.

Support activity for the Master's candidates Giulia Barone, Serena Paesano, Idio Guarino, and Nicola Esposito, and for the scholarship holder Carmen Clemente.

References

[1] A. Dainotti, A. Pescapé, K. C. Claffy, Issues and future directions in traffic classification, IEEE Network 26 (1) (2012) 35–40.

[2] F. Jejdling et al., Ericsson mobility report. Ericsson AB, Business Area Networks, Stockholm, Sweden, Tech. Rep. EAB-18, 4510, 2018.

[3] H. Shi and Y. Li, Discovering periodic patterns for large scale mobile traffic data: Method and applications, IEEE Transactions on Mobile Computing, 2018.

[4] Sandvine, Global Internet Phenomena Spotlight: Encrypted Internet Traffic., 2016.

[5] A. Razaghpanah, A. A. Niaki, N. Vallina-Rodriguez, S. Sundaresan, J. Amann, and P. Gill, Studying TLS usage in Android apps, in 13th ACM CoNEXT, 2017.

[6] B. Saltaformaggio, H. Choi, K. Johnson, Y. Kwon, Q. Zhang, X. Zhang, D. Xu, and J. Qian, Eavesdropping on fine-grained user activities within smartphone apps over encrypted network traffic, in USENIX Workshop on Offensive Technologies (WOOT), 2016.

[7] Y. Fu, H. Xiong, X. Lu, J. Yang, and C. Chen. Service usage classification with encrypted internet traffic in mobile messaging apps, in IEEE Transactions on Mobile Computing, 15(11), 2016.

[8] V. F. Taylor, R. Spolaor, M. Conti, and I. Martinovic, Robust smartphone app identification via encrypted network traffic analysis, in IEEE Trans. Inf. Forensics Security, 13(1), 2018.

[9] G. Aceto, D. Ciuonzo, A. Montieri, and A. Pescapè, Multi-classification approaches for classifying mobile app traffic, Journal of Network and Computer Applications, 103, 2018.

[10] I. Goodfellow, Y. Bengio, and A. Courville, Deep learning. MIT press, 2016.

[11] Z. Wang, The Applications of Deep Learning on Traffic Identification., 2015.

[12] W. Wang, M. Zhu, X. Zeng, X. Ye, and Y. Sheng, Malware traffic classification using convolutional neural network for representation learning, in IEEE International Conference on Information Networking, 2017.

[13] W. Wang, M. Zhu, J. Wang, X. Zeng, and Z. Yang, End-to-end encrypted traffic classification with one-dimensional convolution neural networks, in IEEE International Conference on Intelligence and Security Informatics (ISI), 2017.

[14] M. Lotfollahi, R. Shirali, M. J. Siavoshani, and M. Saberian, Deep packet: a novel approach for encrypted traffic classification using Deep Learning, arXiv, 2017.

[15] M. Lopez-Martin, B. Carro, A. Sanchez-Esguevillas, and J. Lloret, Network traffic classifier with convolutional and recurrent neural networks for Internet of Things, IEEE Access, 5, 2017.

[16] S. E. Oh, S. Sunkam, and N. Hopper, Traffic analysis with deep learning, arXiv, 2017.

[17] J. Ngiam, A. Khosla, M. Kim, J. Nam, H. Lee, and A. Y. Ng, Multimodal deep learning, in 28th International Conference on Machine Learning (ICML), 2011.

[18] P. Syverson, R. Dingledine, and N. Mathewson, Tor: the second generation onion router, in USENIX 13th Security Symposium (SSYM), 2004.

[19] The Invisible Internet Project (I2P), [Online] https://geti2p.net/en/, Jul. 2017.

[20] Project: AN.ON - Anonymity, [Online] http://anon.inf.tu-dresden.de/index_en.html, Jul. 2017.

[21] K. Shahbar and A. N. Zincir-Heywood, Packet momentum for identification of anonymity networks, Journal of Cyber Security and Mobility, vol. 6, no. 1, pp. 27–56, 2017.

[22] Antonio Montieri, Domenico Ciuonzo, Giuseppe Aceto, Antonio Pescapè, Anonymity Services Tor, I2P, JonDonym: Classifying in the Dark (Web), IEEE Transactions on Dependable and Secure Computing (TDSC), Early Access, 2019.